

Flip Distance Between Triangulations of a Simple Polygon is NP-Complete

Oswin Aichholzer*

Wolfgang Mulzer†

Alexander Pilz‡

December 4, 2012

Abstract

Let T be a triangulation of a simple polygon. A *flip* in T is the operation of removing one diagonal of T and adding a different one such that the resulting graph is again a triangulation. The flip distance between two triangulations is the smallest number of flips that is necessary to transform one triangulation into the other.

We show that computing the shortest flip distance between two triangulations of a simple polygon is NP-hard. This complements a recent result that shows NP-hardness of determining the flip distance between two triangulations of a planar point set. For the special case of convex polygons, the problem is equivalent to the rotation distance of binary trees, a central problem which is still open after over 25 years of intensive study.

1 Introduction

Let P be a simple polygon in the plane, that is, the closed region bounded by a piece-wise linear, simple cycle. A *triangulation* T of P is a geometric (straight-line) maximal outerplanar graph whose outer face is the complement of P and whose vertex set are the vertices of P . The edges of T that are not on the outer face are called *diagonals*. Let d be a diagonal whose removal creates a convex quadrilateral f . By replacing d with the other diagonal of f , we again get a triangulation of P . This operation is called a *flip*. The *flip graph* of P is the abstract graph whose vertices are the triangulations of P and in which two triangulations are adjacent if and only if they differ by a single flip. We study the *flip distance*, that is, the minimum number of flips required to transform a given source triangulation into a target triangulation.

Edge flips became popular in the context of Delaunay triangulations. Lawson [8] proved that any triangulation of a planar n -point set can be transformed into any other by $O(n^2)$ flips. Hence, for every planar n -point set the flip graph is connected with diameter $O(n^2)$. Later, he showed that in fact every triangulation can be transformed to the Delaunay triangulation by $O(n^2)$ greedy flips that locally improve the Delaunay property [9]. Hurtado, Noy, and Urrutia [6] gave an example where the flip distance is $\Omega(n^2)$, and they showed that the same bounds hold for triangulations of simple polygons. They also proved that if the polygon has k reflex vertices, then the flip graph has diameter $O(n + k^2)$. This generalizes the well-known fact that the flip distance between any two triangulations of a convex polygon is at most $2n - 10$,

*Institute for Software Technology, Graz University of Technology, Austria. Partially supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18. oaich@ist.tugraz.at.

†Institute of Computer Science, Freie Universität Berlin, Germany. Supported in part by DFG project MU/3501/1. mulzer@inf.fu-berlin.de.

‡Recipient of a DOC-fellowship of the Austrian Academy of Sciences at the Institute for Software Technology, Graz University of Technology, Austria. Part of this work has been done while this author was visiting the Departamento de Matemáticas, Universidad de Alcalá, Spain. apilz@ist.tugraz.at.

for $n > 12$ [14]. The latter case is particularly interesting due to the correspondence between flips in triangulations of convex polygons and rotations in binary trees.

We mention two further remarkable results on the flip graph of triangulations of a point set. Hanke, Ottmann, and Schuierer [5] showed that the flip distance between two triangulations is at most the number of crossings in the overlay of the source and the target triangulation. Eppstein [4] gave a polynomial-time algorithm for computing a lower bound on the flip distance. This bound is tight for point sets that do not contain empty 5-gons; however, except for small instances, such point sets are not in general position (i.e., they must contain collinear triples of points) [1]. For a recent survey on flip operations see Bose and Hurtado [2].

Very recently, the problem of finding the flip distance between two triangulations of a point set was shown to be NP-complete by Lubiw and Pathak [10] and, independently, by Pilz [11]. Here, we show that the corresponding problem for simple polygons is also NP-complete. This result can be seen as a further step towards settling the complexity of deciding the flip distance between triangulations of convex polygons, which is of general interest in theoretical computer science due to its equivalence with the rotation distance between binary trees. This variant of the problem was probably first addressed by Culik and Wood [3] in 1982.

Formally, we study the following problem, called POLY-FLIP: we are given a simple polygon P , two triangulations T_1 and T_2 of P , and an integer l . The question is whether T_1 can be transformed into T_2 by at most l flips. Our reduction proceeds from the RECTILINEAR STEINER ARBORESCENCE PROBLEM, which was shown to be NP-hard by Shi and Su [13]. In Section 2, we describe this problem in more detail, and we give a variant tailored to our needs. Then, in Section 3, we present the *double chains* by Hurtado, Noy, and Urrutia [6], the major building block for our reduction. Finally, in Section 4, we give the details of our hardness proof.

2 The Rectilinear Steiner Arborescence Problem

Our reduction uses a variant of the RECTILINEAR STEINER ARBORESCENCE PROBLEM. Let S be a set of N points in the plane whose coordinates are nonnegative integers. The points in S are called *sinks*. A *rectilinear tree* T is a connected acyclic collection of horizontal and vertical line segments that intersect only at their endpoints. The length of T is the total length of all segments in T (cf. [7, p. 205]). The tree T is a *rectilinear Steiner tree* for S if each sink in S appears as an endpoint of a segment in T . We call T a *rectilinear Steiner arborescence* (RSA) for S if (i) T is rooted at the origin; (ii) each leaf of T lies at a sink in S ; and (iii) for each $s = (x, y) \in S$, the length of the path in T from the origin to s equals $x + y$, i.e., all edges in T point north or east, as seen from the origin [12]. In the RSA problem, we are given a set of sinks S and an integer k . The question is whether S has an RSA of length at most k . Shi and Su showed that the problem is strongly NP-complete; in particular, it remains NP-complete if S is contained in an $n \times n$ grid, with n polynomially bounded in $N = |S|$ [13].¹

We recall an important structural property of the RSA. Let e be a vertical segment of an RSA A for S that does not contain a sink, such that there is a horizontal segment f incident to the upper endpoint a of e . Observe that since A is an arborescence, a has to be the left endpoint of f . Further, suppose that a is not the lower endpoint of another vertical edge. Translate a copy e' of the segment e to the right until e' hits a sink or another segment endpoint of A (which will happen the latest at the right endpoint of f), see Figure 1 for several examples. The segments e and e' define a rectangle R . The upper and the left side of R are completely covered by e and (a part of) f . Since a is the only endpoint incident to both of these sides, all paths from the sinks of A that pass through e or f completely contain these two sides of R , entering the boundary of R at the upper right corner d and leaving it at the lower left corner b . We can therefore reroute every such path at d to continue vertically along the boundary of R

¹Note that a polynomial-time algorithm was claimed [15] that later has been shown to be incorrect [12].

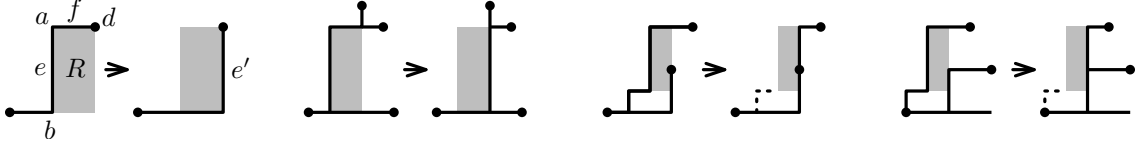


Figure 1: Examples of the slide operation in an RSA. The dots depict sinks and the rectangle R is drawn gray. The dotted segments are removed in the new arborescence, since they do no longer lead to a sink.

until it enters A again (which happens the latest at b) and remove the parts of e and f on R . In the resulting tree we can subsequently remove all segments incident to a leaf that is not a site (which happens if there does not exist a path through b anymore); this gives another RSA A' for S . Observe that A' is not longer than A . We refer to this operation as *sliding e to the right*. The analogous can be done if e is horizontal; e then is *slided upwards*. The *Hanan grid* for a point set P is the set of all vertical and horizontal lines passing through the points in P . In essence, the following theorem can be proven constructively by repeatedly applying slides of segments to any RSA.

Theorem 2.1 ([12]). *Let S be a set of sinks. Then there exists an RSA A for S of minimum length such that all segments of A lie on the Hanan grid for $S \cup \{(0, 0)\}$.* \square

For our reduction we need a restricted version of the RSA problem, called the YRSA problem. In an instance (S, k) of the YRSA problem, we require that no two sinks in S have the same y -coordinate.

Theorem 2.2. *YRSA is strongly NP-complete.*

Proof. Due to the Hanan grid property, the YRSA problem is in NP, as the RSA problem [13]. We show how to reduce RSA to YRSA. Let (S, k) be an instance for an YRSA problem, and label the sinks as $S = \langle s_1, s_2, \dots, s_N \rangle$ in an arbitrary fashion. For $i = 1, \dots, N$, let (x_i, y_i) be the coordinates of s_i and define $s'_i := (x_i N^4, y_i N^4 + i)$. Set $S' := \{s'_1, s'_2, \dots, s'_N\}$. Note that the y -coordinates of the sinks in S' are pairwise distinct.

Now let A be a rectilinear Steiner arborescence for S of length at most a . We can scale A by the factor N^4 and draw a vertical segment from each leaf to the corresponding sink in S' . It follows that there exists an RSA A' of length $a' \leq aN^4 + N^2$.

Suppose there exists an RSA B' of length at most b' . Due to Theorem 2.1, we can assume that B' is on the Hanan grid. We can replace every y -coordinate y_s of every segment endpoint in B' by $\lfloor y_s / N^4 \rfloor N^4$ (ignoring possible segments of length 0). Then this results in an arborescence B for S that was scaled by N^4 (because the resulting drawing remains connected, every path to the origin remains monotone and no cycles are produced since the segments are on the Hanan grid). Any arborescence on the Hanan grid is a union of N paths changing directions at most N times, and every vertical part of such a path is stretched by at most N by the way we changed the y -coordinates. This gives a (very conservative) bound of $bN^4 \leq b' + N^3$ for the length b of B .

Hence, S has an arborescence of length at most k if and only if S' has an arborescence of length at most $kN^4 + N^3$, provided that $N^4 > 2N^3$, that is, $N > 2$. Since the instance $(S', kN^4 + N^3)$ can be computed in polynomial time from (S, k) , and since the coordinates in S' are polynomially bounded in the coordinates of S , it follows that the YRSA problem is strongly NP-complete. \square

3 Double Chains

We use definitions (and illustrations) along the lines of [11]. A *double chain* D consists of two chains, an *upper chain* and a *lower chain*. There are n points on each chain, $\langle u_0, \dots, u_n \rangle$ on

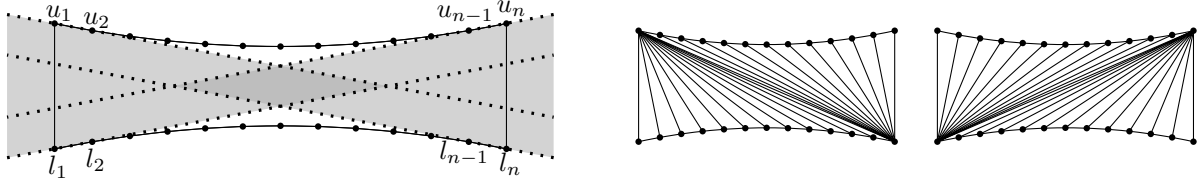


Figure 2: Left: The polygon and the wedge (gray) of a double chain. The diamond-shaped kernel can be stretched arbitrarily by flattening the bend of the chains. Right: The upper and lower extreme triangulations of P_D with a flip distance of $(n-1)^2$, as shown in [6].

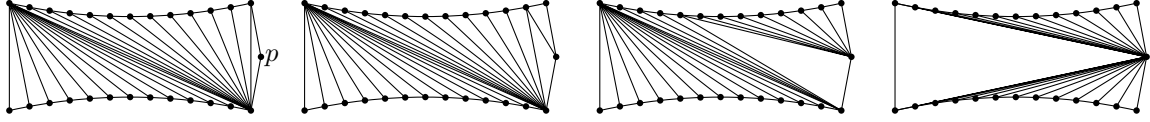


Figure 3: The extra point p in the kernel of D allows flipping one extreme triangulation of P_D^p to the other in $4n-4$ flips.

the upper chain and $\langle l_1, \dots, l_n \rangle$ on the lower chain, both numbered from left to right. The upper chain is reflex w.r.t. any point of the lower chain, and vice versa. Let P_D be the polygon defined by $\langle l_1, \dots, l_n, u_n, \dots, u_1 \rangle$, see Figure 2 (left). We call the triangulation T_u of P_D where u_1 has maximum degree the *upper extreme triangulation*; observe that this triangulation is unique. The triangulation T_l of P_D where l_1 has maximum degree is called the *lower extreme triangulation*. The two extreme triangulations are used to show that the diameter of the flip graph is quadratic, see Figure 2 (right).

Theorem 3.1 ([6]). *The flip distance between T_u and T_l is $(n-1)^2$.* \square

The double chain D is useful for us because by a slight modification of D we can obtain a planar polygon whose flip distance is linear. This introduces a gap in the length of different flip sequences and will allow us in our reduction to enforce a certain structure on short flip sequences. We next define the *kernel* of a double chain.

Definition 3.2. *Let D be a double chain. Let W_1 be the double wedge defined by the supporting lines of u_1u_2 and l_1l_2 whose interior does not contain a point of D , see Figure 2 (left). W_n is defined analogously by the supporting lines of u_nu_{n-1} and l_nl_{n-1} . Let $W = W_1 \cup W_n$ be called the wedge of D . The unbounded set $W \cup P_D$ therefore is defined by four rays and the two chains. A point is outside of D if it is not contained in $W \cup P_D$. The kernel of D is the intersection of the closed half-planes below u_1u_2 and $u_{n-1}u_n$, as well as above l_1l_2 and $l_{n-1}l_n$.²*

Definition 3.3. *Let D be a double chain with its kernel partially to the left of the directed line l_1u_1 , and let p be a point in that part of the kernel. The polygon P_D^p is defined by the sequence $\langle p, l_1, \dots, l_n, u_n, \dots, u_1 \rangle$. The upper extreme and lower extreme triangulations of P_D^p are defined in the same way as for P_D .*

As mentioned in [16], the flip distance between the two extreme triangulations from Figure 2 (right) is much smaller in P_D^p than in P_D . Figure 3 shows how to transform them into each other with $4n-4$ flips. The next lemma shows that this is optimal, even for more general polygons.

Lemma 3.4. *Let P be a polygon that completely contains P_D and has $\langle l_1, \dots, l_n \rangle$ and $\langle u_n, \dots, u_1 \rangle$ as part of its boundary. Further, let T_1 and T_2 be two triangulations that contain the*

²Note that the kernel of D might not be completely inside the polygon P_D (but no point in the kernel is outside D). This is in contrast to the common use of the term “kernel” in visibility problems for polygons.

upper extreme triangulation and the lower extreme triangulation of P_D as a sub-triangulation, respectively. Then T_1 and T_2 have flip distance at least $4n - 4$.

Proof. We follow Lubiw and Pathak [10], slightly generalizing their proof for double chains of constant size.

The triangulation T_1 has $2(n - 1)$ triangles with an edge on one of the chains of D . We refer to the point not incident to that edge as *apex*. For each such triangle at the upper chain, the apex must move from l_n to l_1 , and similarly for the lower chain. We have three types of flips: (1) exchange an edge between the upper and the lower chain by another edge between the two chains; (2) exchange an edge between the two chains by an edge between a point of a chain and a point outside the chain or vice versa; and (3) a flip where less than three of the four points involved are on one of the two chains. A flip of type (1) moves the apex of two triangles by one, a flip of type (2) moves the apex of one triangle from the chain to a point not on the chain or back again, and a flip of type (3) does not move any apex along the chain or between the chain and a point not in the chain. Hence, we can disregard flips of type (3). If moving the apex of a triangle involves at least one flip of type (2), then at least two flips are needed for that triangle to move the apex to a point not on the chain and back again. If moving the apex of a triangle involves no flip of type (2), then at least $n - 1$ flips are needed for that triangle, and each flip moves the apex of one other triangle. We can therefore charge $(n - 1)/2$ flips to each such triangle. Hence, for $n > 5$, the cheapest method is to use flips of type (2), giving the required bound. \square

The following is a special case of a result from [11]. A standalone proof can easily be obtained from the proof of Lemma 3.4 by noting that if the wedge of P_D is empty, there can be no flips of type (2) that move the apex of a triangle from one end of the chain to a different vertex of that chain (as also remarked in [10]).

Lemma 3.5. *Let P be a polygon that completely contains P_D and has $\langle l_1, \dots, l_n \rangle$ and $\langle u_n, \dots, u_1 \rangle$ as part of its boundary, and let T_1 and T_2 be two triangulations that contain the upper extreme triangulation and the lower extreme triangulation of P_D as sub-triangulation, respectively. Suppose there is no vertex in the interior of the wedge of P_D . Then the flip distance between T_1 and T_2 is at least $(n - 1)^2$.* \square

These results on the double chain give the main tools for our reduction. Intuitively, in order to efficiently transform the upper extreme triangulation of a double chain to the lower extreme triangulation, we need a point in the kernel of the double chain and the triangulation of the remaining polygon will have to allow the edges of the extreme triangulation to be flipped to that point.

4 The Reduction

4.1 “Navigating” in a Double Chain

We take a double chain D such that we can add a point z to the right of $l_n u_n$ inside the kernel of D . By P_D^+ , we denote the polygon defined by $\langle l_1, \dots, l_n, z, u_n, \dots, u_1 \rangle$.

Consider a triangulation T of P_D^+ . A *chain edge* is an edge of T between the upper and the lower chain of D . A *chain triangle* is a triangle that contains two chain edges. We use the chain edges to define the *chain path*, an abstract path on the $n \times n$ grid. Let e_1, e_2, \dots, e_m be the chain edges, sorted from left to right according to their intersections with a line ℓ that separates the upper from the lower chain. For $i = 1, \dots, m$, write $e_i = (u_v, l_w)$ and set $c_i = (v, w)$. Note that, in particular, $c_1 = (1, 1)$, which we use as the *root* of our setting. Since T is a triangulation, any two consecutive edges e_i, e_{i+1} share one endpoint, while the other endpoints are adjacent

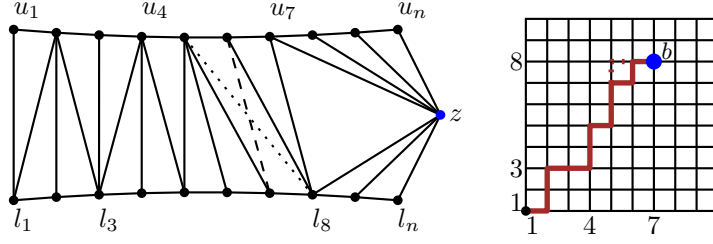


Figure 4: A double chain extended by a vertex z . The vertex z is incident to u_7 and l_8 , represented by the blue point b in the grid. The brown chain path represents the chain triangles. If we flip edges to z , b will move along that path. A flip between chain triangles (dotted edge replaced by the dashed one) changes a bend in that path (from the dotted one).

on the corresponding chain. Thus, c_{i+1} dominates c_i and $\|c_{i+1} - c_i\|_1 = 1$. The chain path is defined as the path $c_1 c_2 \dots c_m$. See Figure 4 for an example.

The chain path is an x - and y -monotone path in the $n \times n$ grid. We call its upper right endpoint b . We now investigate how flipping edges in T affects the chain path.

Observation 4.1. *Suppose we flip an edge that is incident to z . Then the chain path is extended by moving b north or east.* \square

Observation 4.2. *Suppose that T contains at least one chain triangle. When we flip the rightmost chain edge, we shorten the chain path at b .* \square

Finally, we can flip an edge between two chain triangles. This operation is called a *chain flip*.

Observation 4.3. *A chain flip changes a bend from east to north to a bend from north to east, or vice versa.*

Proof. If a chain edge $u_i l_j$ is incident to two chain triangles and is flippable, then the two triangles must be of the form $u_i u_{i-1} l_j$ and $l_j l_{j+1} u_i$, or $u_{i+1} u_i l_j$ and $l_{j-1} l_j u_i$. Thus, flipping $u_i l_j$ corresponds exactly to the claimed change in the chain path. \square

Corollary 4.4. *A chain flip does not change the length of the chain path.* \square

We summarize the results of this section in the following lemma:

Lemma 4.5. *Let T be a triangulation of P_D^+ . Then T uniquely determines an x - and y -monotone path (i.e., the chain path) in the $n \times n$ grid starting at the root $(1,1)$. Conversely, any chain path uniquely determines a triangulation of T . The possible flips of T correspond to the following operations on the chain path: (i) extend the right endpoint north or east; (ii) shorten the path at the right endpoint; (iii) change an east-north bend to an north-east bend, or vice versa.* \square

4.2 Installing Sinks

We show how to reduce YRSA to POLY-FLIP. Let S be a set of N sinks with root at $(1,1)$ on an $(n-1) \times (n-1)$ grid (recall that n is polynomial in N). We describe how to construct a polygon P_D^* for S . Our construction has two integral parameters β and d . With foresight, we set $\beta = 2N$ and $d = nN$.

Let P_D^+ be the polygon from Section 4.1, but with βn vertices on each chain. As we saw in Section 4.1, we can interpret a triangulation of P_D^+ as a chain path in the $\beta n \times \beta n$ grid. We imagine that the sinks of S are in this grid, with their coordinates multiplied by β . For each sink $s = (x, y)$, we place a (rotated) small double chain D_s of size d such that $l_{\beta y}$ and $l_{\beta y+1}$ correspond to the first point on the lower and upper chain of D_s , respectively. In addition,

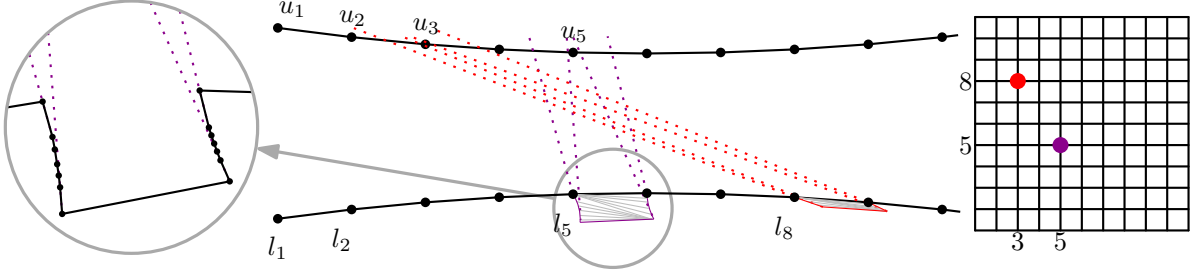


Figure 5: Installing sinks on the grid by constructing the simple polygon P_D^* . The red and the purple double chain can only be “untangled” in reasonable time if they see u_3 and u_5 , respectively (the dotted lines indicate the wedge and the kernel of the corresponding double chain).

$u_{\beta x}$ is the only point in the kernel of D_s and $u_{\beta x}$ is also the only point in the interior of the wedge of D_s , see Figure 5. We call the resulting polygon P_D^* . If β is large enough, the small double chains D_s do not interfere with each other, and P_D^* is simple. Since the y -coordinates in S are pairwise distinct, we create at most one double chain at each edge of the lower chain of P_D^+ . Observe that we have some flexibility for the precise placement of the points of each D_s . Thus we can choose their placement in a way that their coordinates are polynomial in n . See Appendix A for a detailed description of a suitable construction.

Next, we describe the source and target triangulation for P_D^* . The source triangulation T_1 contains all edges of P_D^+ . The interior of P_D^+ is triangulated such that all edges are incident to z , i.e., b is at the root. The small double chains are all triangulated with the upper extreme triangulation. The target triangulation T_2 is defined similarly, but now all the small double chains are triangulated with the lower extreme triangulation (note that the choice of the upper and lower chain is arbitrary for the small double chains).

Hence, each corresponding pair of small double chains in T_1 and T_2 has flip distance $(d-1)^2$ due to Lemma 3.5, unless the appropriate vertex on the upper chain of P_D^* is used. Intuitively, if d is large enough, a shortest flip sequence will have to “traverse” each sink, inducing an arborescence for S . Vice versa, every arborescence for S gives a short flip sequence between T_1 and T_2 . We begin with the simpler direction of the correspondence.

Lemma 4.6. *Let A be an arborescence for S of length k . Then the flip distance on P_D^* between T_1 and T_2 is at most $2\beta k + (4d-2)N$.*

Proof. We describe how to use A to construct a flip sequence of length at most $2\beta k + (4d-2)N$. The triangulations T_1 and T_2 both contain a triangulation T^+ of P_D^+ , such that b is located at the root. We generate flips inside P_D^+ as in Lemma 4.5 that traverse A . This requires $2\beta k$ flips.

Each time b reaches a sink s , we move b north. This creates the appropriate chain triangle for transforming the upper extreme to the lower extreme triangulation of the double chain D_s in $4d-4$ flips, as in Lemma 3.4. Next, we shorten the chain path by moving b back south. This requires two additional flips, so we take $4d-2$ flips per sink. In total, this gives $2\beta k + (4d-2)N$ flips, as claimed. \square

Note that in the above proof, if the traversal of A continues north after visiting a sink s , we can actually save a flip by not shortening the chain path after transforming D_s .

Next we consider the opposite direction of the correspondence. In the proof of the following lemma, we will describe a mapping from each triangulation T of P_D^* to a triangulation T^+ of P_D^+ . For each sink $s \in S$, the corresponding chain triangle t_s in T^+ is defined as the chain triangle in P_D^+ that allows the double chain D_s to be flipped quickly. We say that a flip sequence σ_1 on P_D^+ visits a sink $s \in S$, if σ_1 has at least one triangulation T that contains the corresponding chain triangle t_s . We call σ_1 a *flip traversal* for S if (i) the sequence σ_1 begins and ends in the

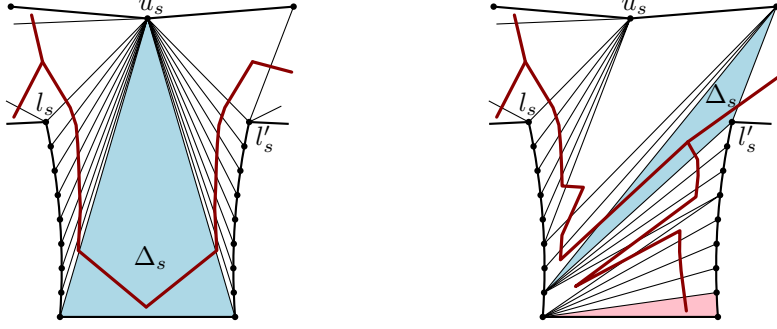


Figure 6: Triangulations of D_s in P_D^* with the triangle Δ_s (blue) incident to the convex vertices of D_s (left) and as an inner triangle when the small chain contains an ear (red). The fat tree indicates the dual.

same triangulation T^+ such that T^+ corresponds to b lying on the root; (ii) the sequence σ_1 visits every sink in S .

Lemma 4.7. *Let σ be a flip sequence on P_D^* from T_1 to T_2 with $|\sigma| < (d-1)^2$. Then there exists a flip sequence σ_1 on P_D^+ such that σ_1 is a flip traversal for S with $|\sigma_1| \leq |\sigma| - (4d-4)N$.*

Proof. Let T^* be a triangulation of P_D^* . An *inner triangle* in T^* is a triangle with only diagonals as sides. An *ear* is a triangle that has two polygon edges as sides. It is well-known that the weak dual of a triangulation of a polygon is a tree where inner triangles have degree 3 and ears correspond to leaves. By construction of P_D^* , the inner triangles of T^* must have a certain form: they have one vertex incident to z (the point in the kernel), or two vertices incident to a small double chain D_s (or both). Observe that in the latter case there can be only one such triangle per D_s .

Let D_s be a small double chain placed between the vertices l_s and l'_s with u_s being the vertex in the kernel of D_s . We define Δ_s as the triangle that is either the inner triangle incident to two vertices of D_s or the triangle that is incident to both convex vertices of D_s but is not an ear. Note that in the first case the third vertex might be u_s and that in the latter case the third vertex has to be u_s . Due to the structure of P_D^* there always exists exactly one such triangle Δ_s per sink: if the convex vertices of D_s are not part of an ear, then Δ_s is the triangle between them and u_s . Otherwise we follow the path in the weak dual from the ear in D_s ; if the next triangle is not an inner one, it has to have an edge of the small double chain as a side, but there is only a limited number of such edges. Note that Δ_s might be $l_s l'_s u_s$. See Figure 6.

For each sink s , let the polygon $P_{D_s}^{u_s}$ consist of the double chain D_s extended by the vertex u_s (recall Definition 3.3), and let T_s denote a triangulation of it. We define a mapping of any triangulation T^* of P_D^* to a triangulation T^+ of P_D^+ and to triangulations T_s for all sinks s . The triangulation T^+ contains every triangle that has all three vertices in P_D^+ . For each triangle ∇ that has two vertices on P_D^+ and one on the left chain of D_s , we replace the apex on D_s by l_s . The analogous is done if the apex of a triangle ∇ is on the right chain of D_s ; we replace that apex by l'_s . For every sink s , the triangle Δ_s is known to have an apex at a point u_i of the upper chain. In T^+ , we replace Δ_s by the triangle $l_s l'_s u_i$. Since these are exactly the triangles needed for a triangulation of P_D^+ and no two triangles overlap, T^+ is indeed a triangulation of P_D^+ . Similarly, all triangles in T^* that have all three vertices among the ones of $P_{D_s}^{u_s}$ are also in T_s , and the triangles having two points on D_s and whose apex is not in $P_{D_s}^{u_s}$ get their apex at u_s in T_s (note that this includes Δ_s). See Figure 7.

Now we show that a flip in T^* corresponds to at most one flip either in T^+ or in precisely one T_s for some sink s . We do this by considering all the possibilities for two triangles that share a common flippable edge. Note that by construction no two triangles mapped to triangulations of different polygons $P_{D_s}^{u_s}$ and $P_{D_t}^{u_t}$ can share an edge (with $t \neq s$ being another sink).

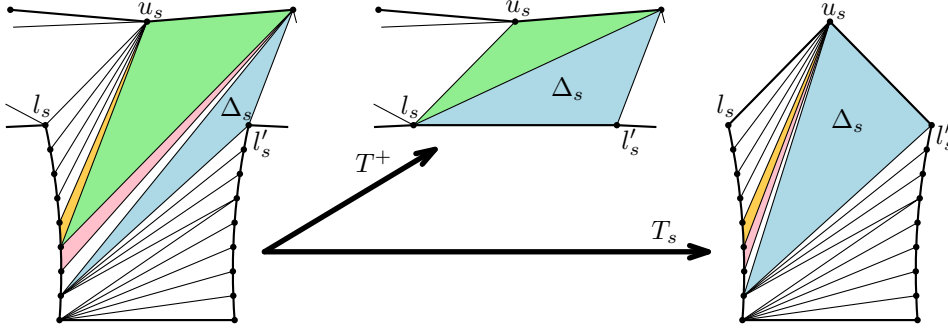


Figure 7: A part of a triangulation of P_D^* and the two corresponding triangulations T^+ and T_s .

Case 1. We flip an edge between two triangles that are either both mapped to T^+ or to T_s and are different from Δ_s . This flip clearly happens in at most one triangulation.

Case 2. We flip an edge between a triangle Δ that is mapped to T_s and a triangle ∇ that is mapped to T^+ , such that both Δ and ∇ are different from Δ_s . This results in a triangle Δ' that is incident to the same edge of $P_{D_s}^{u_s}$ as Δ , and a triangle ∇' having the same vertices of P_D^+ as ∇ . Since the apex of Δ' is a vertex of the upper chain or z (otherwise, it would not share an edge with ∇), it is mapped to u_s , as the apex of Δ . Also, the apex of ∇' is on the same chain of D_s as the one of ∇ . Hence, the flip affects neither T^+ nor T_s .

Case 3. We flip the edge between a triangle ∇ mapped to T^+ and Δ_s . By construction, this can only happen if Δ_s is an inner triangle. The flip affects only T^+ , because the new inner triangle Δ'_s is mapped to the same triangle in T_s as Δ_s , since both apices are moved to u_s .

Case 4. We flip the edge between a triangle Δ of T_s and Δ_s . Similar to Case 3, this affects only T_s , because the new triangle Δ'_s is mapped to the same triangle in T^+ as Δ_s , since the two corners are always mapped to l_s and l'_s .

By the choice of d and Lemma 3.5, all triangulations T_s have to be transformed by flipping edges to u_s . Thus, the triangle $l_s l'_s u_s$ must occur in some triangulation in σ_1 . For each s , we need $4d - 4$ flips to transform T_s , and no flip changes more than one triangulation we mapped T^* to. The lemma follows. \square

4.3 Arborescences and Traces

Next, we define traces (domains drawn on the grid) and relate them to flip traversals. A *trace* is drawn on the $\beta n \times \beta n$ grid. It consists of *edges* and *boxes*: an edge is a line segment of length 1 whose endpoints have positive integer coordinates; a box is a square of side length 1 whose corners have positive integer coordinates. Similar to arborescences, we require that a trace R (i) is (topologically) connected; (ii) contains the root $(1, 1)$; and (iii) from every grid point contained in R there exists an x - and y -monotone path to the root that lies completely in R . We say R is a *covering trace* for S (or, R *covers* S) if every sink in S is part of R .

Let σ_1 be a flip traversal as in Lemma 4.7. By Lemma 4.5, we can interpret the sequence σ_1 as the evolution of a chain path. This gives a covering trace R for S in the following way. For every flip in σ_1 that extends the chain path, we add the corresponding edge to R . For every chain flip in σ_1 , we add the corresponding box to R . Afterwards, we remove from R all edges that coincide with a side of a box in R . Clearly, R is (topologically) connected. Since σ_1 is a flip traversal for S , every sink is covered by R (i.e., incident to a box or edge in R). Note that every grid point p in R is connected to the root by an x - and y -monotone path on R , since at some point p belonged to a chain path in σ_1 . Hence, R is indeed a trace, the unique *trace* of σ_1 .

Next, we define the *cost* of a trace R , $\text{cost}(R)$, so that if R is the trace of a flip traversal σ_1 , then $\text{cost}(R)$ gives a lower bound on $|\sigma_1|$. An edge has cost 2. Let B be a box in R . A *boundary side* of B is a side that is not part of another box. The cost of B is 1 plus the number of

boundary sides of B . Then, $\text{cost}(R)$ is the total cost over all boxes and edges in R . E.g., the cost of a tree is twice the number of its edges, and the cost of an $a \times b$ rectangle is $ab + 2(a + b)$. An edge can be interpreted as a degenerated box, having two boundary sides and no interior.

Proposition 4.8. *Let σ_1 be a flip traversal and R a trace for σ_1 . Then $\text{cost}(R) \leq |\sigma_1|$.*

Proof. We argue that every element of R has unique corresponding flips in σ_1 that account for its cost. Let e be an edge of R . Then e corresponds to at least two flips in σ_1 : one that extends the chain path to create e , and one that removes e (because the chain path starts and ends in a single point). Next let B be a box in R . The interior of B corresponds to at least one chain flip in σ_1 . Moreover, when adding the box for a chain flip to the trace, we either transform edges to boundary sides or make boundary sides disappear from the boundary of the new trace. See Figure 8 for examples. However, when a chain flip adds a new box B to a trace, B is adjacent to at least two already existing elements (edges or boundary sides). Hence, by induction, the new boundary edges of a box add at most the cost that the box removes. \square

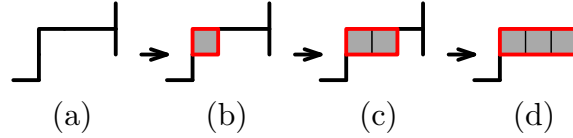


Figure 8: Examples of how boundary sides (red) are added to a trace. To a trace of cost 16 (a) a box (gray) is added (b), which transforms two edges in boundary sides and adds two boundary sides, resulting in an overall cost of 17. The next box removes one boundary side and one edge and adds three boundary sides (c), the cost becomes 18. A box might also remove more than two elements (d), reducing the overall cost to 17.

Now we relate the length of an arborescence for S to the cost of a covering trace for S , and thus to the length of a flip traversal. Since each sink (s_x, s_y) is connected in R to the root by a path of length $s_x + s_y$, traces can be regarded as generalized arborescences. In particular, we make the following observation.

Observation 4.9. *Any shortest path tree A_{σ_1} in R for the root w.r.t. S is an arborescence.* \square

If σ_1 contains no chain flips, the corresponding trace R has no boxes, but it may not be acyclic. However, due to Observation 4.9 it contains an arborescence A_{σ_1} , in particular with $2|A_{\sigma_1}| \leq \text{cost}(R)$. The next lemma extends this relation between the length of an arborescence for S and the cost of a covering trace for S (and therefore the length of a flip traversal of S) to the case where chain flips occurred in σ_1 .

Lemma 4.10. *Let σ_1 be a flip traversal of S . Then there exists a covering trace R for S in the $\beta n \times \beta n$ grid such that R does not contain a box and such that $\text{cost}(R) \leq |\sigma_1|$.*

Proof. There exists at least one trace of cost at most $|\sigma_1|$, namely the trace of σ_1 . Let \mathcal{R}_1 be the set of all covering traces for S that have minimum cost. If \mathcal{R}_1 contains a trace without boxes, we are done. Otherwise, every covering trace in \mathcal{R}_1 contains at least one box.

Let $\mathcal{R}_2 \subseteq \mathcal{R}_1$ be those covering traces among \mathcal{R}_1 that contain the minimum number of boxes. Let $Q \in \mathcal{R}_2$, and let B be a *maximal* box in Q , i.e., Q has no other box whose lower left corner has both x - and y -coordinate at least as large as the lower left corner of B . We investigate the structure of Q . Note that the property of being a trace is invariant under mirroring the plane along the line $x = y$; in particular, the choice of B in Q as a maximal box remains valid.

Observation 4.11. *Every corner c of B is incident either to a sink, an edge, or another box.*

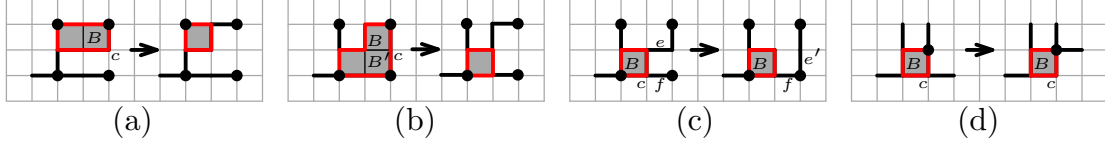


Figure 9: Parts of traces to be modified; the boundary sides are shown in red. (a) A box that has a corner c with no incident elements can be removed. (b) Two adjacent boxes that have a shared corner c without any incident elements can be removed. (c) Replacing a single edge. (d) Sliding an edge.

Proof. If not, we could remove c and B while keeping the sides of B not incident to c as edges, if necessary, see Figure 9(a). The resulting structure would be a trace with smaller cost, contradicting the choice of Q . \triangle

Observation 4.12. *Suppose B shares a horizontal side with another box B' . Let c be the right endpoint of the common side. Then c is incident either to a sink, an edge, or another box.*

Proof. Suppose this is not the case. Then we could remove B and B' from Q while keeping the sides not incident to c as edges, if necessary, see Figure 9(b). This results in a valid trace that has no higher cost but less boxes than Q , contradicting the choice of Q . \triangle

Observation 4.13. *Let c be the lower right corner of B . Then c has no incident vertical edge.*

Proof. Such an edge would be redundant, since c already has an x - and y -monotone path to the root that goes through the lower left corner of B . \triangle

Now we derive a contradiction from the choice of Q and the maximal box B . Note that since β is even, all sinks in S have even x - and y -coordinates. We distinguish two cases.

Case 1. There exists a maximal box B whose top right corner c' does not have both coordinates even. Suppose that the x -coordinate of c' is odd (otherwise, mirror the plane at the line $x = y$ to swap the x - and the y -axis). By Observation 4.11, there is at least one edge incident to the top right corner of B (it cannot be a box by the choice of B , and it cannot be a sink because of the current case). Recall the slide operation for an edge in an arborescence. This operation can easily be adapted in an analogous way to traces. If there is a vertical edge v incident to c' , it cannot be incident to a sink. Thus, we could slide v to the right (together with all other vertical edges that are above v and on the supporting line of v). Hence, we may assume that c' is incident to a single horizontal edge e ; see Figure 9(c). By Observation 4.11, the bottom right corner c of B must be incident to an element. We know that c cannot be the top right corner of another box (Observation 4.12), nor can it be incident to a vertical segment (Observation 4.13). Thus, c is incident to an element f that is either a horizontal edge or a box with top left corner c . But then e could be replaced by a vertical segment e' incident to f , and afterwards B could be removed as in the proof of Observation 4.11, contradicting the choice of Q .

Case 2. The top right corner of each maximal box has even coordinates. Let B be the rightmost maximal box. As before, let c be the bottom right corner of B . The y -coordinate of c is odd; see Figure 9(d). By the choice of B , we know that c is not the top left corner of another box: this would imply that there is another maximal box to the right of B . We may assume that c is not incident to a horizontal edge, as we could slide such an edge up, as in Case 1. Furthermore, c cannot be incident to a vertical edge (Observation 4.13), nor be the top right corner of another box (Observation 4.12). Thus, B violates Observation 4.11, and Case 2 also leads to a contradiction.

Thus, the choice of Q forces a contradiction in either case. No trace of minimum cost contains a box, so every minimum trace is an arborescence. This completes the proof of Lemma 4.10. \square

Corollary 4.14. *Let σ be a flip sequence on P_D^* from T_1 to T_2 with $|\sigma| < (d-1)^2$. Then there exists an arborescence A_σ in the $\beta n \times \beta n$ grid for S with $2|A_\sigma| \leq |\sigma| - (4d-4)N$.*

Proof. Use Lemma 4.7 to obtain a sequence σ_1 . Then apply Lemma 4.10 to the trace of σ_1 . \square

Corollary 4.15. *Let σ be a flip sequence on P_D^* from T_1 to T_2 with $|\sigma| \leq 2\beta k + (4d-2)N$. Then there exists a rectilinear Steiner arborescence for S of length at most k .*

Proof. Trivially, there always exists an arborescence on S of length less than $2nN$, so we may assume that $k < 2nN$. Hence (recall that $\beta = 2N$ and $d = nN$),

$$2\beta k + 4dN - 2N < 2 \times 2N \times 2nN + 4nN^2 - 2N < 12nN^2 < (d-1)^2,$$

for $n \geq 14$ and positive N . We can therefore apply Corollary 4.14, which gives an arborescence A that covers every sink and has length at most $\beta k + N$. Hence there is an arborescence A' for S that is not longer than A and that is on the Hanan grid. Therefore, the length of A' is a multiple of β . Thus, since $\beta > N$, we get that A' has length at most βk , so the corresponding arborescence for S on the original grid has length at most k . \square

Now we have all the tools to prove our main theorem.

Theorem 4.16. *POLY-FLIP is NP-complete.*

Proof. Since the flip distance in polygons is polynomially bounded, the problem is in NP. We reduce YRSA to POLY-FLIP. Let (S, k) be an instance of YRSA such that S lies on a grid of polynomial size. Construct a polygon P_D^* and triangulations T_1 and T_2 as described at the beginning of Section 4.2. This takes polynomial time. Set $l := 2\beta k + (4d-2)N$. By Lemma 4.6 and Corollary 4.15, the sinks S allow a tour of length at most k if and only if T_1 and T_2 have flip distance at most l . \square

References

- [1] Z. Abel, B. Ballinger, P. Bose, S. Collette, V. Dujmović, F. Hurtado, S. Kominers, S. Langerman, A. Pór, and D. Wood. Every large point set contains many collinear points or an empty pentagon. *Graphs Combin.*, 27:47–60, 2011.
- [2] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.
- [3] K. Culik II and D. Wood. A note on some tree similarity measures. *Inf. Process. Lett.*, 15(1):39–42, 1982.
- [4] D. Eppstein. Happy endings for flip graphs. *JoCG*, 1(1):3–28, 2010.
- [5] S. Hanke, T. Ottmann, and S. Schuierer. The edge-flipping distance of triangulations. *J.UCS*, 2(8):570–579, 1996.
- [6] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22:333–346, 1999.
- [7] F. Hwang, D. Richards, and P. Winter. *The Steiner Tree Problem*. Annals of Discrete Mathematics. North-Holland, 1992.
- [8] C. L. Lawson. Transforming triangulations. *Discrete Math.*, 3(4):365–372, 1972.
- [9] C. L. Lawson. Software for C^1 surface interpolation. In J. R. Rice, editor, *Mathematical Software III*, pages 161–194. Academic Press, NY, 1977.

- [10] A. Lubiw and V. Pathak. Flip distance between two triangulations of a point-set is NP-complete. In *Proc. 24th Canadian Conference on Computational Geometry*, pages 127–132, Charlottetown, Canada, August 2012.
- [11] A. Pilz. Flip distance between triangulations of a planar point set is NP-complete. *ArXiv e-prints*, June 2012, 1206.3179.
- [12] S. K. Rao, P. Sadayappan, F. K. Hwang, and P. W. Shor. The rectilinear Steiner arborescence problem. *Algorithmica*, 7:277–288, 1992.
- [13] W. Shi and C. Su. The rectilinear steiner arborescence problem is np-complete. In *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 780–787, 2000.
- [14] D. Sleator, R. Tarjan, and W. Thurston. Rotation distance, triangulations and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–682, 1988.
- [15] V. Trubin. Subclass of the Steiner problems on a plane with rectilinear metric. *Cybernetics*, 21:320–324, 1985.
- [16] J. Urrutia. Algunos problemas abiertos. In N. Coll and J. Sellares, editors, *Proc. IX Encuentros de Geometría Computacional*, pages 13–24. Univ. De Girona, July 2001.

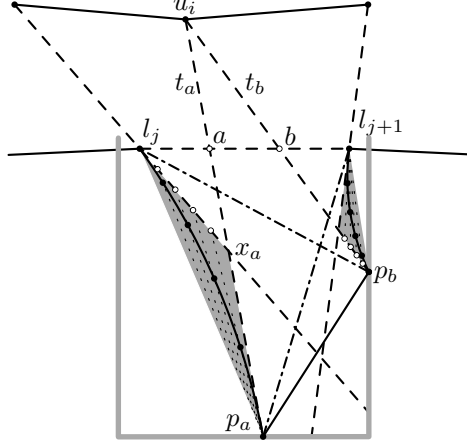


Figure 10: Construction of a small double chain for a sink.

A A Note on Coordinate Representation

Since it is necessary for the validity of the proof that the input polygon can be represented in size that is bounded by a function polynomial in the size of the YRSA instance, we give a possible method on how to embed the polygon with vertices at rational coordinates whose numerator and denominator are polynomial in N .

One can use the well-known parametrization of the unit circle, choosing n points with rational coordinates inside the wedge defined by the two common tangents of two instances of the unit circle for the upper and lower chain. Given these points, we now construct the small double chains for each sink. See Figure 10. Recall that, since β is a multiple of two, there are no small double chains on neighboring positions on the lower chain. Hence, for each sink we w.l.o.g. can define an orthogonal region within which we can safely draw the small double chain; we call this region the *bin* of the sink (outlined gray in Figure 10). Consider a sink (i, j) . We first partition the segment $l_j l_{j+1}$ into thirds to obtain two points a and b , which again have rational coordinates; note that these points are not part of the polygon but “helper points” for our construction. Let t_a and t_b be the lines through u_i and a , and through u_i and b , respectively. The lines t_a and t_b intersect the bin at the points p_a and p_b , respectively. These two points will be the endpoints of the two chains. In addition to t_a the supporting lines of $u_{i-1} l_j$, as well as $l_j p_b$ and $p_a l_{j+1}$ define the triangular region $l_j p_a x_b$ (shaded gray in Figure 10) wherein we may place the chain incident to l_j (note that, e.g., only one of t_a and $p_a l_{j+1}$ will bound the triangular region). The chain incident to l_{j+1} is constrained analogously. We place a circular arc C through p_a and l_j inside $l_j p_a x_a$. C can be chosen to be tangent to either $p_a x_a$ or $x l_j$, in order to be contained in $l_j p_a x_a$. It is well-known that, for a line with rational slope through a point with rational coordinates on C that intersects C in a second point, this second point has rational coordinates as well. Suppose that C is tangent to $p_a x_a$. Then we divide the line segment $l_j x_a$ into $(d - 1)$ parts (where d is the number of elements on a small chain). A line through a point defined by this subdivision and p_a gives a rational point on C . Likewise, we can choose the points if C is tangent to $l_j x_a$. The points for the second chain are chosen analogously.

The coordinates are rational, and since every point can be constructed using only a constant number of other points, the numerator and denominator of each point are polynomial.